

-2-

Remarks/Arguments

As noted at page 1, lines 12-30 of the Applicants' specification, most large businesses operate LANs at several sites and interconnect these LANs using dedicated circuits leased from Network Service Providers. As explained in the Applicants' specification, the leased dedicated circuits are sized to meet the maximum bandwidth requirements of the interconnections and frequently operate below capacity making inefficient use of the Network Service Provider network.

Many of the LANs operated by large businesses operate according to IEEE 802.1 standards. These standards provide protocols, among them 802.1q, that enable the businesses to partition their LANs into multiple Virtual LANS (VLANs). So, for example, a large business may partition its LANs into separate VLANs for different departments or operations of the business, like Finance, Manufacturing, Design and Legal.

The Applicants' networks, routing devices and packet routing methods go beyond the VLAN capabilities required by IEEE 802.1q to enable Network Service Providers to provide interconnection of a very large number of VLANs for a large number of distinct customers on shared Service Provider network facilities in a manner which makes more efficient use of the Service Provider network than the typical leased dedicated circuits, while preserving isolation between the data communications of separate customers, and while preserving the VLAN partitioning of those customers within their own data networks. Such shared Service Provider network facilities can comprise data switches or routing devices at multiple nodes distributed across one or more extended metropolitan areas, and transmission facilities linking the data switches or routing devices as shown, for example, in Figure 1 of the Applicants' specification.

As noted at page 1, line 31 to page 2, line 11 of the Applicants' specification, the 12 bit capacity of the VLAN tag specified by the IEEE 802.1q standard limits the number of distinct VLANs to 4095. Network Service Providers need to support many more than 4095 distinct customers on a shared network. Moreover, many customers of the Network Service Providers are already using the IEEE 802.1q VLAN identifier to partition their own networks and do not want Network Service Providers to disrupt such partitioning by changing the VLAN identifiers on packets traversing the Service Provider networks or by restricting their use of particular VLAN identifiers. The Applicants' invention does not rely on IEEE 802.1q VLAN identifiers to isolate VLAN from VLAN and customer from customer.

-3-

Consequently, the Applicants' invention is not limited by the restricted size of the IEEE 802.1q VLAN identifier space, and permits each customer to assign VLAN Identifiers without considering how other customers may be assigning VLAN Identifiers.

The Applicants claim communications networks, routing devices for such communications networks and methods of routing packets through such communications networks in which each packet entering a network at an ingress virtual port is assigned a respective *egress address* and routed through the communications network according to that respective egress address. The assigned egress address corresponds to a respective destination address of the entering packet when a correspondence between the destination address and an egress address is known. When no correspondence between the destination address and an egress address is known, the assigned address is a *broadcast egress address* which is *selected based on the ingress virtual port* to correspond to a distinct set of virtual ports, the distinct set of virtual ports comprising the ingress virtual port. Consequently, any broadcasting of the entering packet is restricted to the distinct set of virtual ports that includes the ingress virtual port.

The claimed networks, routing devices and packet routing methods enable the Applicants to interconnect large numbers of customer LAN segments for large numbers of different customers without relying on VLAN identifiers to isolate the VLANs and the customers from one another. Because the VLAN identifiers are not relied upon for routing of the packets, the number of customer VLANs that the Applicants can interconnect is not limited by the restricted number of valid IEEE 802.1 VLAN identifiers, and the VLANs and customers are isolated even if different customers use the same VLAN identifiers in their LAN segments.

The Examiner rejected claims 1-7, 9-17, 19-27, 29-30, 32-34, 41-42 and 44-56 under 35 USC 103(a) as being unpatentable over Ross (US 5,394,402) in view of Stone (US 6,041,057).

In the Applicants' claims 1-34 and 41-56, an *"egress address"* is assigned to each packet entering a *connectionless* network via an ingress virtual port. The egress address corresponds to a respective destination address of the entering packet when a correspondence between the destination address and an egress address is known. The egress address is a *broadcast egress address selected based on the ingress virtual port* to correspond to a *unique set of virtual ports* to which the ingress virtual port belongs when no correspondence between the destination address and an egress address is known. The assigned egress address is added to the packet, and the packet is routed across the network according to the egress address to the egress virtual port

-4-

without establishing a connection across the network for routing the packet. The egress address is removed from the packet at the egress port.

Ross discloses LANs interconnected across a backbone network to form Virtual LANs or (VLANs). Packets entering the interconnecting network from a LAN segment connected to an "internal port" of a hub are assigned *VLAN designations* which are used to route the packet to other internal ports of the hub or to an "external port" for transmission to other hubs over the backbone network. Ross discloses that the backbone network may be a packet network or an ATM network. Ross provides no details of the addressing applied at the external port to route the packet between hubs, but teaches that the *VLAN designations* are used to route the packets to internal ports at the destination hubs for egress from the interconnecting network.

The Examiner argued that Ross discloses "distinct sets" of virtual ports "where no virtual port belongs to more than one of the distinct sets". However, column 3, lines 35-36 say that "any one of the internal ports may be assigned more than one VLAN designation", so if the VLAN designations define Ross's sets of internal ports, the sets are overlapping sets, not *distinct sets* as recited in Applicants' claims. Without distinct sets of internal ports, Ross could not guarantee secure separation of traffic from different Service Provider customers as required by one of the Applicants' intended applications.

The Examiner further argued that Ross assigns an "egress address" to each packet entering the interconnecting network. While Ross provides no details of the "start and addressing for backbone network" shown as item 92 in Figure 6, it is clear from column 8, lines 60-66 that this addressing merely routes the packet from the "external port" of one hub to the "external port" of a destination hub. The packet is routed within the hub to the "internal port" *at which it egresses the interconnecting network* based on the VLAN designation or the destination address or both (column 10, lines 13-19), not based on the "addressing for the backbone network". Consequently, the "addressing for the backbone network" is not an *egress address* as recited in Applicants' claims.

Because Ross uses standard VLAN designations to isolate VLANs, Ross does not address Service Provider issues described at page 1, line 31 to page 2, line 11 of Applicants' specification. In particular Ross is limited by the finite number of VLAN designations permitted by the 12 bit standard VLAN designation, and cannot cope with use of the same VLAN designations by different Service Provider customers. Consequently Ross does not address important

-5-

practical problems that can be addressed by Applicants' claimed invention. In practical terms, Ross is really only applicable within the network of an Enterprise. Ross does not address the problems of a Service Provider trying to provide VLAN service to many distinct customers over shared Service Provider network.

The Examiner stated that Ross fails to disclose the use of a broadcast address when no egress address is known for a destination address. However, the Examiner stated that Stone discloses "mapping of an unknown destination address to a broadcast VLAN address", and that it would be obvious to a skilled person to "map an unknown destination address to a broadcast VLAN address" in the network of Ross "to determine VLAN connectivity".

The Applicants note that Stone nowhere mentions a "broadcast VLAN address" and the Applicants are not aware of such terminology being used in the relevant field. Moreover, the Applicants claims and specification do not recite a "broadcast VLAN address" and the Applicants are not aware of any such concept.

Stone discloses VLANs interconnected across an ATM network using ATM tag switching. Because the ATM network is *connection-oriented*, Stone requires a protocol whereby edge switches of the ATM network initiate tagged point-to-point and point-to-multipoint virtual connections for forwarding end-user messages between VLAN segments across the ATM network (column 2, line 38 to column 3, line 43).

The ATM switches forward the end-user messages by associating VLAN identifiers in the end-user messages with tag values of the established tagged virtual connections. The tag values are used to route the end-user messages link by link via the tagged virtual connections. At each successive switch in the tagged virtual connections, the tag values are translated into tag values for a next link of the virtual connection until the end-user messages reach destination edge switches of the ATM network (column 2, line 44 to column 3, line 10).

VLAN packets comprising the end-user messages are segmented into fixed length ATM cells at the edge switch where they enter the ATM network (column 12, lines 49-51), and must be reconstructed at the edge switch where they leave the ATM network.

Stone initiates a point to multipoint tagged virtual connection for each VLAN supported by his network (column 6, lines 38-42). Tag values are assigned to "unknown unicast messages"

-6-

and the "unknown unicast messages" are forwarded along the appropriate point to multipoint virtual connections according to the assigned tags (column 7, lines 46-51).

Consequently Stone teaches a *connection-oriented* network in which packets are routed through the network according to *identifiers associated with pre-established connections* unlike the Applicants' *connectionless* network in which packets are routed according to *assigned egress addresses*. Unlike the Applicants, Stone requires considerable overhead to set up and maintain a large number of virtual *connections* through the network. Stone requires further additional processing because, unlike the Applicants' *assigned egress addresses*, the *ATM tags* used by Stone to route packets through his *connection-oriented network* simply provide the routing to the next switch in the ATM network and are changed at each transit switch — i.e. the ATM tags of Stone do not point to the egress point of the network at which the packet leaves the network and are not "*egress addresses*". The Applicants' "*egress addresses*" can be used to route packets at multiple nodes across the network without translation because they identify one or more egress ports of the network. Consequently, the Applicants' claimed approach is very different in practical terms from the Stone's disclosure because Stone requires significantly more processing at transit nodes to carry a packet across the ATM network than does the Applicants' claimed invention to carry the packet across Applicants' network. Moreover, Stone incurs further processing overhead by breaking incoming LAN frames into ATM cells of standard length (column 12, lines 49-51), and by therefore needing to reconstruct the LAN frames at the egress nodes where the end-user messages are transferred to other VLAN segments.

There is no teaching or suggestion in either Ross or Stone that would motivate a skilled reader to combine these references, and the Examiner has arrived at the particular hypothetical combination advanced in his arguments based only on the hindsight afforded by the Applicants' specification.

Moreover, like Ross and unlike the Applicants' claims, Stone *relies on VLAN identifiers* to define VLAN membership and to isolate VLANs. Column 5, lines 32-34 state that Stone's switches share VLAN identifiers in topology advertising messages. Column 5, lines 56-61 and column 6, lines 19-21 state that Stone's switches learn VLAN membership from the advertised VLAN identifiers. Column 6, lines 38-53 describe the set-up of point-to-multipoint tagged virtual connections for each VLAN based on the learned VLAN membership. Column 12, lines 26-48 describe use of the point-to-multipoint tagged virtual connections to route "unknown unicast end-

-7-

user messages" within a particular VLAN. Consequently, the point-to-multipoint routing of "unknown unicast messages" depends on VLAN membership *determined from VLAN identifiers*. Because different customers could use the same VLAN identifiers, Stone (like Ross) would not provide the customer isolation required in a Carrier Service Provider network application. Even if Stone were to require that different customers use different VLAN identifiers to preserve customer isolation, a requirement that would not be acceptable to many customers, Stone's network (like Ross's) could still only support 4095 distinct VLANs. This limitation would not meet the needs of many Service Provider carrier network applications. The Applicants overcome these limitations by *not* relying on IEEE 802.1 VLAN identifiers set by Carrier Service Provider customers in their customer networks.

Furthermore, a skilled reader seeking to combine Stone with Ross would note that Ross discloses an ATM embodiment (see Ross Figure 7 and column 7, lines 58-66). As Stone teaches an ATM network using tagged virtual connections, the most straightforward combination of Stone with Ross would build on Ross's ATM embodiment resulting in a *connection-oriented* ATM network using tagged virtual connections unlike the Applicants' *connectionless* network which, as discussed above, avoids the significant overhead processing needed set up and maintain virtual connections. Any other hypothetical combination of Ross and Stone could only be guided by the hindsight afforded by the Applicants' specification and claims.

-8-

Conclusion

In summary, there is no motivation to combine Ross and Stone in the manner suggested by the Examiner, and any such combination would necessarily *rely on VLAN Identifiers* to isolate different VLANs, a strategy which does not provide reliable isolation between customers when VLANs of multiple customers are interconnected without restricting the customers' use of the VLAN identifiers. The Applicants' reliance on distinct sets of virtual ports and *assigned egress addresses* to achieve isolation between VLANs and customers avoids this problem. Moreover, the most straightforward combination of Ross and Stone would result in a *connection-oriented* network in which packets are routed according to *Identifiers of pre-established connections*, and such a *connection-oriented* network would require considerably more overhead processing to set up and maintain the pre-established connections than the Applicants' *connectionless* network which routes packets according to *assigned egress addresses* and does not require the processing overhead needed for pre-establishment and management of connections through the network. Applicants therefore submit that claims 1-7, 9-17, 19-27, 29-30, 32-34, 41-42 and 44-56 are patentable over Ross and Stone, and request reconsideration and allowance of these claims.

The Examiner stated that claims 8, 18, 28, 31 and 35-40 are allowed. The Applicants acknowledge the allowance of these claims with thanks.

The Examiner stated that claim 43 is objected to as being dependent on a rejected base claim, but would be allowable if rewritten in independent form. The Applicants thank the Examiner, but submit that claim 43 is allowable in its present form for the reasons stated above in the discussion of the rejected base claim.

-9-

The Applicants request reconsideration and allowance of the application in view of the arguments detailed above.

Respectfully submitted,

DELANEY, David MacDonald, et al.


By: C.W. Junkin
Registration No. 32,312

CWJ/lw

c/o NORTEL NETWORKS LIMITED
IP Law Group
P.O. Box 3511, Station C
Ottawa, Ontario, Canada K1Y 4H7
Telephone: (613) 768-3013
Date: October 19, 2004